GEORASTER: Commercial Database Infrastructure for Geospatial Imagery

Jim Farley

Imagery and image-based applications present both unique opportunities and specialized challenges to producers, publishers, and consumers of these digital data products. The structure of image data, the frequency of acquisition, and the demand for products with increasingly higher resolution all result in larger and larger data sets. Often time becomes a multiplier as well as the same portion of the earth is imaged over and over to develop a time series for evaluation and analysis to understand changing conditions on the ground. Frequently image data must exist in a highly secure, controlled environment to support emergency response and other secure applications. Finally, many applications using imagery have a collaborative component in a setting that requires many users to access, use and in some cases modify these data.

Support for large data sets, highly scaled environments, and stringent security requirements supporting a broad community of users with differing roles all suggest that commercial database technology should play a key role for producers and consumers of image data. Historically this has not been possible because commercial DBMS products provided no support for georeferenced imagery. However, the recently released Oracle 10g database supports image-types in the database, manages virtually unlimited data volumes and provide features such

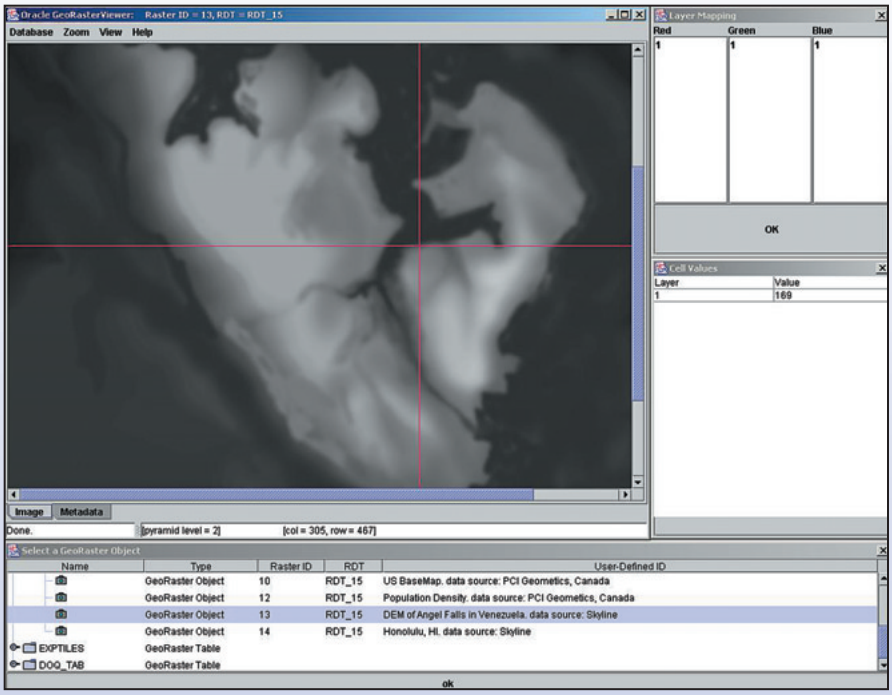

Figure 1 *Oracle10g can store and manage raster imagery and grid-based data for both Enterprise GIS and core business applications.*

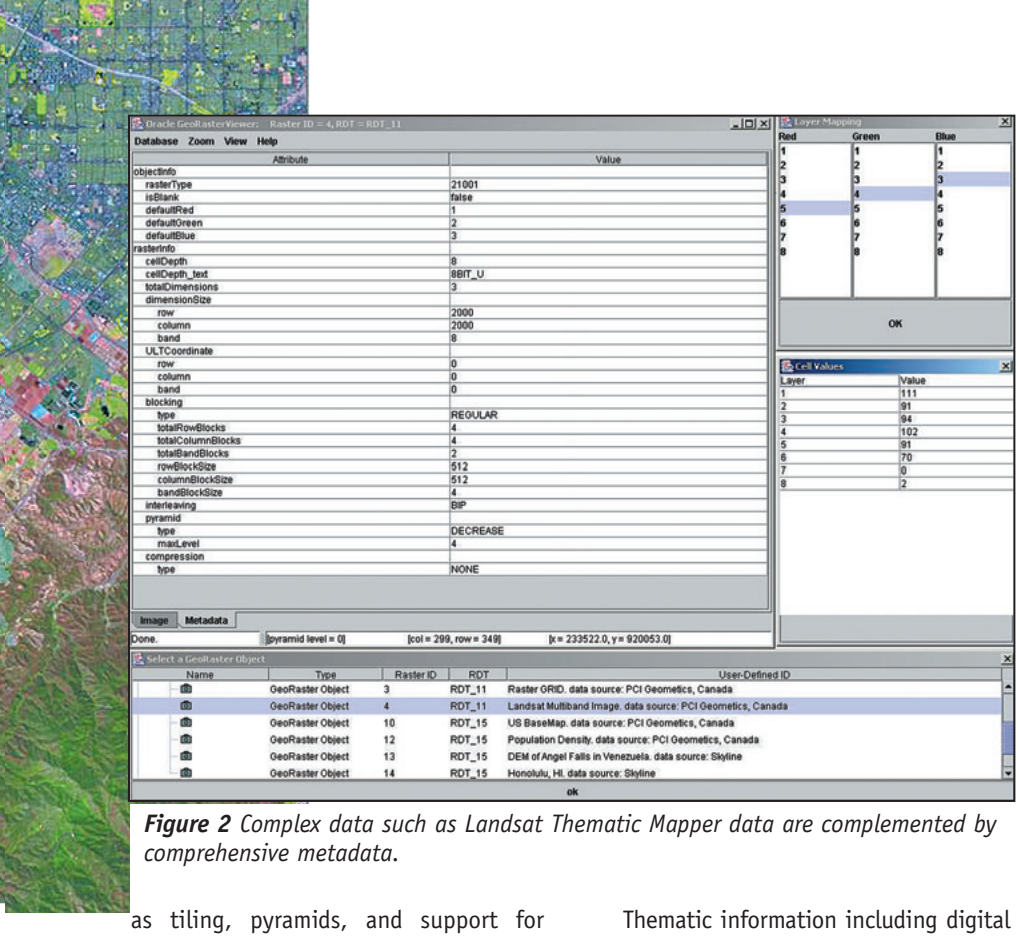

Figure 2 *Complex data such as Landsat Thematic Mapper data are complemented by comprehensive metadata.*

as tiling, pyramids, and support for multi-band data. These features make this database platform a practical option for imagery users and producers.

Geographic Raster Data Management

The use of raster data has evolved from classified applications in defense and intelligence over the past fifty years to become a common tool in planning, security and surveillance, business intelligence, agriculture, and in a wide range of transportation and natural resource applications. In addition to these traditional consumers of raster data, the insurance industry and real estate are increasing their use of georeferenced raster data while, entertainment, media and medical imaging incorporate these data using local or on-earth based coordinate systems.

There are two basic kinds of raster data considered here: grid-based data and raster data:

- Grid data typically has a uniform grid in which each grid-cell has specific attribute values measuring things such as elevation, frequency, concentration, etc. Given information about the bounding coordinates of the grid, the location of each cell can be calculated. Thematic information including digital terrain elevation data, land use information, pollution concentration, land cover information, geological information, and rainfall information can be effectively used in this model.
- Digital imagery or pixels-based data created from optical or other sensors is collected using a variety of technologies including satellite remote sensing, airborne photogrammetry, and sonar. Digital orthophoto images and images composed of one or more bands measuring reflectance along the electromagnetic spectrum are popular for documenting the land base at a specific point in time; as a base for feature extraction; or as a means of exploring complex relationships in the physical environment.

As we will see in the discussion below both grid-based and raster imagery are managed in the Oracle 10g database.

That's a Lot of Data

One key benefit associated with imagery and remotely sensed raster data, the frequency with which these data are or can be acquired, has also been a barrier to its broad adoption in many sectors: the sheer volume of data that results. Airborne platforms can be scheduled with relative ease and moderate expense and a typical satellite platform orbits the earth multiple times a day. Because the satellite will photograph the same portion of the earth at regular, predictable intervals it is the most cost-effective way to track changes on the ground. However, the resulting raster datasets tend to be large, even when compression is a realistic option for a specific application. Individual images are large and they may accumulate at a rapid rate. For instance, the IKONOS platform from Space Imaging orbits the earth once every 98 minutes or 14 times a day and in a relatively short time has produced more than 100 million square kilometers of imagery.

Clearly, such image-based sensor devices and the data they collect offer great opportunities to extend the application contexts in which these data are used. It is equally clear that these platforms produce large quantities of data that can impose significant data management, manipulation and storage requirements on the consumer. The Oracle Spatial 10g GeoRaster feature is specifically engineered to meet these requirements and overcome many of the barriers that may have hindered broader adoption of these valuable tools.

Oracle 10g GeoRaster

GeoRaster is a feature of Oracle 10g Spatial that supports storage, indexing, querying, manipulation; analysis, and the delivery of image-based and gridded raster data and its associated metadata. This feature has been developed to deliver enterprise-class data management capability to organizations that have large-scale image processing requirements. GeoRaster is a new spatial data type that includes an object-relational schema; comprehensive metadata; base methods to manipulate raster data for management purposes; and other infrastructure support functions such as load and export. Analysis functionality for image processing, interpretation, change detection, and domain decision support applications, etc. is provided by industry partners that are leaders in remote sensing, photogrammetry, and geospatial applications.

The GeoRaster data and schema objects are used to store multidimensional gridded data and raster layers

that are referenced to the Earth's surface or to a local coordinate system. Pixel-level management, retrieval, and analysis are supported. GeoRaster is designed to meet the general needs of broad application groups including:

- Traditional GIS and remote sensing applications—users manage their geographic raster and gridded data assets using a scaleable, secure, and robust RDBMS for defense, intelligence, agriculture, natural resource management.
- Business applications—leverage raster-based data in conjunction with other basic location data (address, etc.) to inventory and evaluate site locations and to track fixed and/or continuous assets. Asset management and facilities management applications in energy and the utilities are likely consumers.
- Image and Gridded Raster Data Repositories/Clearinghouses—clearinghouse servers that need to ingest, store, and disseminate very large volumes of geoimagery benefit from the scalability and the security.

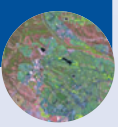

Benefits of Managing Raster Data in Oracle 10g

By integrating the complete range of information (attribute, georeferenced raster and spatial vector data) in a single environment the complexity of individual work flows in an enterprise is simplified and thereby less likely to be affected by error as a by-product of merging and exchanging data between "special-purpose" systems. Further, this means that the cost of building and fielding applications that use this data will be reduced while the security, scalability and reliability of the application environment will improve dramatically. In addition, the following benefits are derived from this integrated approach:

- Seamless geographic coverage possible, even for extensive geographic coverages at fine-grained resolution, with scalable infrastructure capable of supporting exabyte requirements.
- Common semantics (SQL) across all enterprise data.
- Support by leading 3rd party image processing, GIS and visualization tools.
- All data in this secure infrastructure can be published via Web services as needed.

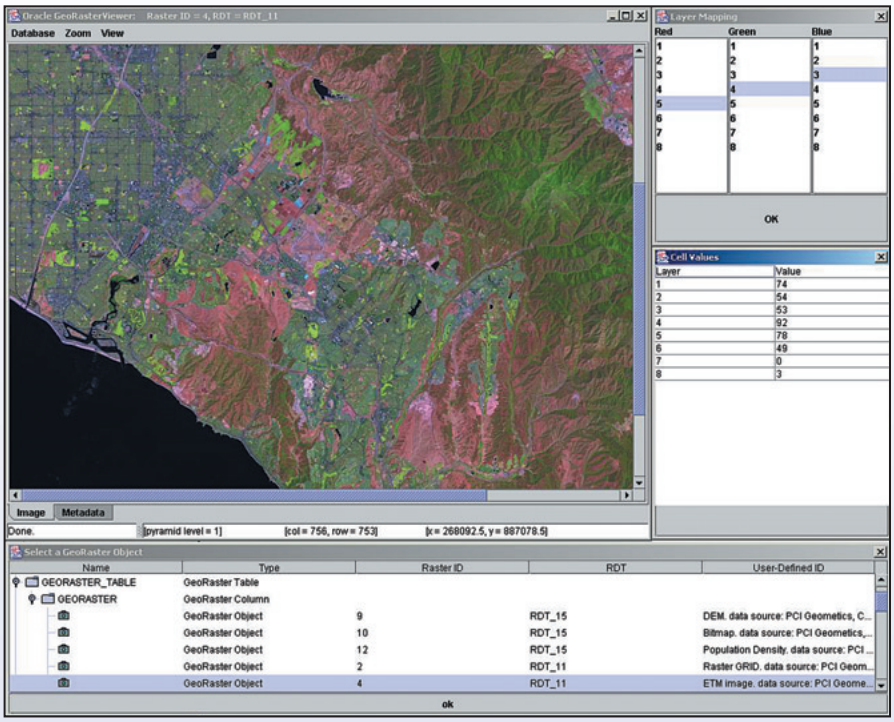

Figure 3 *Using high-quality raster data products provides infrastructure to Core business applications such as Field Service, Asset Management, and Supply Chain.*

The Application of GeoRaster

There is a wide range of use cases for GeoRaster. The section below highlights use cases for deploying the Oracle Spatial 10g GeoRaster feature with third party analysis and visualization tools.

Defense and Security

- Problem—To detect threats to security and to ascertain threat levels while managing an appropriate response matrix both friendly and unfriendly assets need to be tracked.
- Context—Detecting changes over time on the ground is a prerequisite to effective planning and response.
- Solution—Aerial and satellite remote sensing platforms are tasked to generate images from an area of interest on a regular basis. Images are collected, georeferenced, and loaded into a database as raster or gridded layers. Client tools are used to examine current images in association with historic images for the same surface coverage archived in the database. The current raster is mosaicked with surrounding images to create seamless coverage of area of interest.
- Critical Fact—Image data is the only geographic information that can be acquired in a scheduled/timely manner (tasked) for a specific local providing rapid access to current, accurate geodata.

Emergency Response

- Problem—Assess on-the-ground damage and develop suitable response scenario given damage to critical infrastructure.
- Context—Making a rapid assessment of damage to infrastructure in the aftermath of a disaster event is critical to a timely response and mitigation.
- Solution—Raster data from aerial and satellite remote sensing platforms are used to compare "before and after" conditions on the ground. GeoRaster data is used in conjunction with associated vector data for continuous asset infrastructure (*e.g.*, road, rail, power grid, gas, telco) to determine a) damage b) response scenario c)viable corridors for first responders and d) routing to appropriate facilities.
- Critical Fact—Only raster data can provide the near real-time data acquisition needed to accomplish damage assessment and meet the time critical requirements of first responders. A raster data layer provides an ideal

backdrop to display infrastructure data (*e.g.*, pipelines, transmission lines) in a readily comprehensible form.

Enterprise Asset Management

- Problem—Optimize preventative maintenance, field service and operations across a network of continuous and fixed assets (*e.g.*, stations, substations and pipeline).
- Context—Accomplish ongoing monitoring and management across a variety of assets to support normal preventative maintenance and operations.
- Solution—Incorporate raster data from aerial photography into the existing geospatial data used to map assets under one single enterprise database management system. Raster data used in conjunction with vector information representing property boundaries, lease zones and easements enable field service personnel to save time and wear and tear on equipment, zeroing in on problem areas on the ground.
- Critical Fact—Raster data increases the efficiency of resources in the field and drives down costs making operations more efficient.

State and Local Government

- Problem—State and local government have limited resources to manage zoning, tax assessment, etc.
- Context—State and local government budgets are stretched like no time in recent memory. With tax base often shrinking and public service expenses increasing, local officials are compelled to improve efficiency and to maximize coordination and communication between local departments in government.
- Solution—a base map of digital raster data (digital orthophoto quads or contracted photogrammetric product) when stored in the enterprise database in state government provides a common frame of reference that can be used across multiple departments (Transportation, Tax Assessment, Zoning, School Administration, etc.) to support decision making and the formulation and implementation of policy.
- Critical Fact—Raster data provides a common frame of reference that can be used across departments and state agencies.

Agriculture Monitoring

- Problem—National agricultural agencies need to document and verify agricultural utilization of the land.
- Context—Farmers often are required to report the land use practices they employ and the crops they seed allowing officials to project earnings based on yield. In addition, it is common under some circumstances to compensate a farmer for land that is left fallow or unseeded.
- Solution—Use of aerial image data acquired on an annual basis in conjunction with vector information from the land management agency enable officials and local agriculturalists to create an accurate record of acres in crop and acres fallow.
- Critical Fact—Only image-based information acquired on an annual basis can provide the baseline information needed to make the assessments needed to support equitable taxation and remuneration.

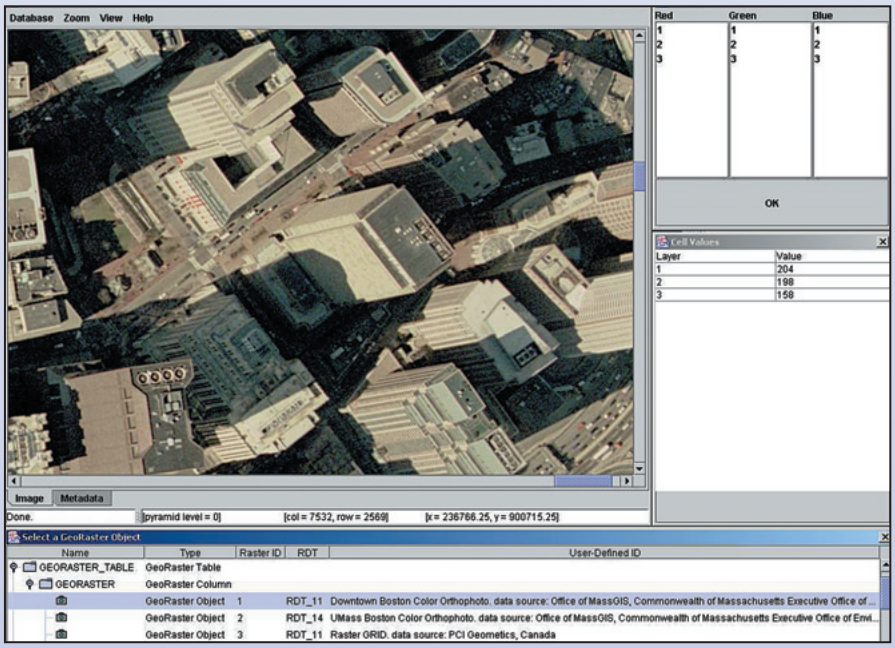

Figure 4 *Raster and grid-based data stored in GeoRaster will support complex visualizations and animated fly-overs using partner solutions.*

Summary

Raster imagery and grid-based data are key ingredients for enterprise applications that leverage, or might leverage, geospatial data to address core business requirements or to improve the way in which strategic business decisions are formulated. Historically, these data have been isolated from the majority of the core enterprise data because the commercial database management systems used to store core business data could not accommodate these rich image-based data products in any meaningful way. This in turn limited the degree to which these data were applied in business and the range of applications in which it could be effectively applied. With the recent release of the Oracle 10g database platform these barriers to the effective, integrated use of imagery and grid-based data have been removed. Oracle Spatial 10g is uniquely positioned to deliver enterprise-class support for storing and managing raster data in the context of a high performance, scalable, secure environment. By taking this valuable raster-based spatial data out of the file system, where it exists in an insecure, often transient state and storing it as named types in the secure environment of the world's leading database management system.

About the Author

Jim Farley leads Oracle Spatial Product initiatives in raster technology, hosted location-based services, and in the integration of location technologies across Oracle's eBusiness Suite Applications.